REMARKS

Marked-up pages showing changes made to the disclosure and claims are attached hereto.

Care has been taken to ensure that no new matter has been added by this amendment.

Receipt of an initial Office Action on the merits is awaited. Should there be any questions or concerns regarding this amendment, the Examiner is invited to contact the Applicant's undersigned representative by telephone.

The Commissioner is hereby authorized to charge the required fees to Deposit Account No. 19-5113.

Respectfully submitted,
MAGED E. BESHAI

By:_____
Max R. Wood, Reg. No. 40,388
Agent of Record

Date: Jan 17th, 2003

ADDRESS: OGILVY RENAULT
1981 McGill College Avenue
Suite 1600
Montreal, Quebec, Canada
H3A 2Y3

VERSION WITH MARKINGS TO SHOW CHANGES MADE

DISCLOSURE:

Paragraph from Page 2, line 25 to Page 3, line 5 (Amended)

This is accomplished in a rate-controlled multi-class high-capacity packet switch described in Applicant's copending United States Patent Application No. 09/244,~~284~~ 824 which was filed on February 4, 1999. Although the switch described in this patent application is adapted to switch variable sized packets at very high speeds while providing grade-of-service and quality-of-service control, there still exists a need for a distributed switch that can form the core of a powerful high-capacity, high-performance network that is adapted to provide wide geographical coverage with end-to-end capacity that scales to hundreds of Tera bits per second (Tbs), while providing grade of service and quality of service controls.

Paragraphs from Page 4, line 31 to Page 7, line 27 (Amended)

~~It is an object of the invention to provide a very high-capacity packet switch with a distributed core that is adapted to provide guaranteed quality of service, as well as providing intra-switch data paths with a granularity that reduces or eliminates a requirement for tandem switching.~~

~~The invention therefore provides a high-capacity packet switch that includes a plurality of core modules that operate in a circuit switching mode, and a plurality of edge modules that are connected to subtending packet sources and subtending packet sinks, each of the edge modules operating in a packet switching mode. The core modules~~

switch payload traffic between the edge modules using wavelength division multiplexing (WDM) and time division multiplexing (TDM).

Each of the core modules is preferably a space switch. Any of the well known textbook designs for a space switch can be used. However, the preferred space switch is an electronic single stage rotator switch, because of its simple architecture, ease of control and scalability. A one of the edge modules is preferably co-located with each core module and serves as a controller for the core module.

Each of the edge modules has a plurality of ingress ports and a plurality of egress ports. Each of the ingress ports has an associated ingress queue. An ingress scheduler sorts packets arriving in the ingress queues from the subtending packet sources, the sort being by destination edge module from which the respective packets are to egress from the high capacity packet switch for delivery to the subtending packet sinks. The ingress scheduler periodically determines a number of packets waiting in the ingress queues for each other respective edge module, and sends a capacity request vector to each of the controllers of the core modules. The capacity request vector indicates a current required capacity from the sending ingress edge module to one or more egress edge modules. The capacity request vector sent to a given controller relates only to a group of channels that connect the edge module to the given core module.

Each edge module also maintains a vector of pointers to the sorted payload packets, the vector of pointers being arranged in egress edge module order. A scheduling matrix having one row corresponding to each time slot of a time frame and each egress edge module is associated with the vector of pointers and determines a data transfer schedule for the ingress edge module.

Each ingress edge module also maintains an array of reconfiguration timing circuits, a one of the reconfiguration timing

circuits being associated with each of the core modules. The reconfiguration timing circuits are respectively coordinated with time clocks in the respective edge modules that serve as controllers for the core modules. The reconfiguration timing circuits enable reconfiguration of channel switching in the core modules using a short guard time.

Each core module preferably comprises a plurality of single-stage rotator switches. Each rotator switch preferably accommodates a number of input channels equal to the number of ingress edge modules, as well as a number of output channels equal to the number of egress edge modules. In a folded edge module configuration, each edge module preferably has one channel connected to an input port and one channel connected to an output port of each single stage rotator switch. In an unfolded edge module configuration, each edge module is either an ingress module or an egress module. The ingress and egress modules are preferably arranged in co-located pairs. In the unfolded configuration, each ingress edge module preferably has one channel connected to an input port of each rotator switch. Each egress module likewise preferably has one channel connected to an output port of each rotator switch.

The invention also provides a method of switching payload data packets through a distributed data packet switch. In accordance with the method, payload data packets are received from a subtending source at an ingress edge module of the distributed data packets switch. An identity of an egress edge module from which the data packets should egress from the distributed data packet switch is then determined. Using the identity of the egress edge module, the data packets are arranged in a sorted order with other data packets received so that the data packets are in a sorted order corresponding to the identity of the edge module from which the data packet should egress from the

~~distributed data packet switch. The sorted data packets are transferred in fixed length data blocks that are switched through the core module to the egress module. The fixed length data blocks contain concatenated packets of variable length, and the respective egress module parses the variable size packets according to methods known in the art. The fixed length data blocks are switched through the core module to the egress module. Thereafter, the payload data packet is transferred to a subtending sink.~~

A very high-capacity packet switch is adapted to provide a high service-quality, as well as providing intra-switch data paths with a fine granularity that reduces or eliminates a requirement for tandem switching. The packet switch requires a scheduler to coordinate the transfer of packets across the switch, and the scalability of the switch is primarily determined by the throughput of its scheduler. Providing a fine granularity in a high-capacity switch requires an extensive scheduling effort that may not be realizable with a single controller. The invention, therefore, provides a switch that includes a plurality of core modules that operate in a time-division mode, and a plurality of edge modules that are connected to subtending packet sources and sinks, with each core module having its own controller which includes a packet-transfer scheduler.

In accordance with an aspect of the present invention, there is provided a packet switch. The packet switch comprises a plurality of independently-controlled core modules, a plurality of ingress modules, and a plurality of egress modules. The packet switch may further include a plurality of core controllers operating concurrently and independently; one core controller associated with each of the independently-controlled core modules and having a packet scheduler. Each ingress module receives packets from subtending packet sources and has a link directed to each of the core modules. Each egress

module has a link from each of the core modules and transmits packets to subtending packet sinks. Each of the ingress modules is operable to issue packet-transfer requests and distribute the packet-transfer requests among the core modules for scheduling. Each core module computes schedules in response to receiving packet-transfer requests, the schedules specifying time slots in a predefined time frame for each request. The ingress modules and the core modules can be geographically distributed and each ingress module is provided with a plurality of timing circuits each communicating with a time counter associated with one of the core modules to realize time coordination between each ingress module and the core modules.

In accordance with another aspect of the present invention, there is provided a method of scheduling. The method is performed by a controller of a core module having $S \geq 1$ space switches and at least one link to each of a plurality of ingress modules, where each ingress module formulates capacity-allocation requests preferably organized in capacity-request vectors each entry of which specifying an input port p, an output port π, and a number K of time slots per time frame. The method relies on a data structure to facilitate the scheduling process. The data structure preferably comprises a first three-dimensional matrix having a space dimension s representing space switches associated with the core module, a space dimension p representing space-switch input ports, and a time dimension t representing the time slots in a slotted frame, and a second three-dimensional matrix having the space dimension s, a space dimension π representing space-switch output ports, and said time dimension t. The method comprises steps of creating the data structure, receiving capacity-allocation requests from the ingress edge modules, selecting a space switch s and a time slot t and, if both entries {s, p, t} of the first three-dimensional matrix and {s,π, t} of the second three-dimensional matrix are free, allocating

the space switch s and the time slot t and marking entries {s, p, t} and {s,π, t} as busy. The step of selecting is repeated until at most K time slots are allocated. The method includes the further step of terminating a current connection by setting the value of K to equal to zero.

In accordance with a further aspect of the present invention, there is provided a distributed packet switch. The distributed packet switch comprises a plurality of m cross connectors, a plurality of n core modules, a plurality of m×n edge modules, and a plurality of n core controllers each having a core scheduler. Each cross connector has n outer links and n inner links. Each outer link connects to an edge module and includes Λ channels in each direction to and from the edge module. Each inner link connects to a core module and includes Λ channels in each direction to and from the core module. Each core module comprises a number of space switches not exceeding the ratio Λ/n. The edge modules and the core modules can be spatially distributed over a wide geographical area and the outer and inner links are preferably wavelength-division-multiplexed links. Each edge module has means for time coordination with the core modules. The core controller of any core module is adapted to compute a schedule in response to receiving capacity-allocation requests, the schedule specifying, for each capacity-allocation request, time slots in a predefined time frame.

In accordance with a still further aspect of the present invention, there is provided a packet switch. The packet switch comprises a plurality of egress modules, each for transmitting packets on at least one network link, a plurality of ingress modules, each for receiving packets from at least one network link and capable of requesting ingress-to-egress-module connections for transferring

received packets to any other of the egress modules, and a plurality of core modules, each capable of simultaneously receiving and independently responding to the ingress-to-egress-module connection requests from any of the ingress modules and of providing the ingress-to-egress module connections between any of the ingress modules and any of the egress modules in response to the connection requests. Each core module may have its own controller for allocating and scheduling resources to the ingress-to-egress-module connections. A core controller operates independently of, and concurrently with, the other core modules' controllers. Each edge module has a plurality of ingress ports each having an associated ingress buffer for receiving packets from subtending packet sources. An ingress controller in each ingress module sorts packets arriving in the ingress buffer into ingress queues, each ingress queue corresponding to an egress module from which packets are to egress from the switch for delivery to subtending packet sinks. Each edge module has a number of timing circuits at least equal to the number of core modules, each of the timing circuits being time-coordinated with a time counter associated with each of the core modules.

In accordance with an additional aspect of the present invention, there is provided a method of switching packets through a switch having a plurality of ingress modules each having at least one ingress port, a plurality of egress modules each having at least one egress port, and a plurality of core modules. Each ingress module is coupled to each core module, each core module is coupled to each egress module, and a packet can traverse only one ingress module, one core module, and one egress module in moving from an ingress port to an egress port. The method comprises steps of receiving, at an ingress module, packets from subtending traffic sources, the ingress module selecting the egress modules to which to send the packets, sending

connection requests to selected core modules, and requesting connections of specified capacities. The method includes the further steps of determining a feasible capacity allocation in response to a connection request, subtracting the feasible capacity allocation from a specified capacity, and returning an updated connection request to the ingress module that issued the connection request. If the feasible capacity allocation is less than the specified capacity, the ingress module may send the connection request to another core module.

In accordance with another aspect of the present invention, there is provided an ingress module in a packet switch. The ingress module comprises an ingress controller, a plurality of ingress ports each having an ingress buffer for receiving packets from subtending packet sources where each packet indicates one of predefined destinations, a plurality of output ports for directing the packets to a plurality of core modules, means for sorting the packets received in the ingress buffer into ingress queues each corresponding to one of the destinations, means for storing a set of predefined paths to each of the predefined destinations, means for formulating connection requests, each connection request specifying a destination and a required capacity allocation, and means for selecting a candidate path from among the predefined paths for each connection request.

In accordance with a further aspect of the present invention, there is provided a core module in a packet switch. The core module comprises at least one space switch having a plurality of input ports and a plurality of output ports, and a core controller adapted to receive connection requests, each connection request specifying a required capacity allocation and a destination selected from among a set of predefined destinations. The core controller provides means for associating each destination with one of the output ports, and a scheduler associated with the core controller times the transfer of packets from the

input ports to the output port and communicates scheduling results to sources of the connection requests.

Paragraphs from Page 21, line 27 to Page 22, line 32 (Amended)

The capacity-request matrix 44 sent to a core module 34 is normally a sparse matrix with a majority of null entries since the capacity demand is split among eight core modules. The controller for a core module attempts to schedule the capacity requested by each ingress edge module 22 using data structures generally indicated by references 46 and 48. Each of the data structures 46, 48 is a three-dimensional matrix having a first space dimension *s*, which represents the respective space switches associated with the core module 34; a second space dimension *p*, which represents the space-switch ports; and a time dimension *t*, which represents the slots in a slotted frame. Thus, an entry in data structure 46 is represented as *{s,p,t}*. The second dimension *p* may represent an input channel, if associated with the data structure 46, or an output channel if associated with the data structure 48. If the number of slots *T* per frame is 16, for example, then in the configuration of FIG. 1, which shows a centralized core, the size of the three-dimensional structure 46 is $128 \times 256 \times 16$. In the distributed core shown in FIG. 3, each core module uses a three-dimensional structure 46 of size $16 \times 256 \times 16$.

When the connections through a core module 34 are reconfigured, the core controller may either reschedule the entire capacity of the respective core module 34 or schedule capacity changes by simply ~~perturbating~~ perturbing a current schedule. If the entire capacity of the core module is reconfigured, each ingress edge module 22 must communicate a complete capacity request vector to the core module while, in the latter case, each ingress edge module 22 need only report capacity request changes, whether positive or negative, to a

respective core controller. A negative change represents capacity release while a positive change indicates a request for additional capacity. The incremental change method reduces the number of steps required to prepare for reconfiguration.

Paragraph from Page 23, line 21 to Page 24, line 2 (Amended)

A capacity request is rejected by a core module if sufficient matching slots cannot be found. In order to reduce the incidence of mismatch, the matching process should always start from a selected space switch at a selected time slot and follow the same search path for each capacity request. For example, the matching process may start from space switch 0 at time slot 0 and then proceed by increasing the time slot, s, from 0 to *ST*, where ~~*S*~~ *T* is the number of time slots per ~~channel~~timeframe. It then continues to the next time-port plane 53 until the 16 planes (in this example) are exhausted or the capacity is successfully allocated, whichever takes place first. The result produced by this packing search, which is well known in the art, is an occupancy pattern shown in FIG. 8.

CLAIMS:

Claims 1-14 cancelled.

~~1. A high capacity packet switch, comprising:~~

~~a plurality of core modules, each of the core modules operating in a circuit switching mode;~~

~~a plurality of edge modules connected to subtending packet sources and subtending packet sinks, each of the edge modules operating in a packet switching mode;~~

~~wherein the core modules switch payload traffic between the edge modules using wavelength division multiplexing (WDM) and time division multiplexing (TDM).~~

2. A high capacity packet switch as claimed in claim 1 wherein each core module is a space switch.

3. A high capacity packet switch as claimed in claim 2 wherein each core module is a single stage electronic rotator switch.

4. A high capacity packet switch as claimed in claim 1 wherein:

each edge module has a plurality of ingress ports, each of the ingress ports having an associated ingress queue; and

an ingress scheduler sorts packets arriving in the ingress queues from the subtending packet sources, the sort being by egress edge module from which the respective packets are to egress from the high capacity packet switch for delivery to the subtending packet sinks.

5. A high capacity packet switch as claimed in claim 4 wherein the ingress scheduler periodically determines a number of packets waiting in the ingress queues for each respective egress edge module and sends a capacity request vector to each of the controllers of the core modules, the capacity request vector sent to a given controller relating only to a group of channels that connect the edge module to the given core module.

6. A high capacity packet switch as claimed in claim 5 wherein each ingress edge module maintains a vector of pointers to the packets sorted by egress edge module and a scheduling matrix that provides a port number for each slot in which a data block can be transferred, the scheduling matrix being arranged in the same egress edge module order so that the scheduling matrix and the pointers are

~~logically aligned; and, when a non-blank entry in the scheduling matrix indicates an egress port through which a data block can be transferred, a corresponding pointer in the vector of pointers is used to locate a starting point for the data block in the packets waiting in the ingress queues.~~

~~7. A high capacity packet switch as claimed in claim 1 wherein the core modules and the edge modules are spatially distributed.~~

~~8. A high capacity packet switch as claimed in claim 7 wherein one edge module is co-located with each core module, and the edge module serves as a controller for the core module.~~

~~9. A high capacity packet switch as claimed in claim 8 wherein each edge module has *M* reconfiguration timing circuits, where *M* is the number of core modules, each of the reconfiguration timing circuits being time-coordinated with a time counter in the respective edge modules that serve as processors for the core modules, to coordinate data transfer from the ingress edge modules when the core modules are reconfigured to change channel connectivity.~~

~~10. A high capacity packet switch as claimed in claim 1 wherein each edge module is connected to each core module by at least one communications link.~~

~~11. A high capacity packet switch as claimed in claim 10 wherein each core module comprises a plurality of single-stage rotator switches, each rotator switch having a number of input ports collectively adapted to accommodate a number of channels equal to the number of ingress edge modules and a number of output ports collectively adapted to accommodate a number of channels equal to the number of egress~~

edge modules, and each edge module has at least one channel to each of the rotator switches.

12. A high capacity distributed packet switch, comprising:

a plurality of distributed core modules, each core module having a plurality of input ports and a plurality of output ports, the distributed core modules switching payload traffic in a circuit switching mode; and

a plurality of distributed ingress edge modules, each ingress edge module having a plurality of ingress ports for receiving payload traffic from subtending sources and a plurality of egress ports for transferring payload traffic to the core modules;

a plurality of egress edge modules having a plurality of ingress ports for receiving payload traffic from the core modules and a plurality of egress ports for transferring the payload traffic to subtending sinks; and

each of the ingress and egress edge modules operates in a packet switching mode.

13. A high capacity distributed packet switch as claimed in claim 12 wherein the ingress edge modules and the egress edge modules comprise integrated units of one ingress edge module and one egress edge module each.

14. A method of switching payload data packets through a distributed data packet switch, comprising the steps of:

a) receiving a payload data packet from a subtending source at an ingress edge module of the distributed data packets switch;

b) determining an identity of an egress edge module from which the data packet should egress from the distributed data packet switch;

~~c) arranging the data packet in a sorted order with other data packets received so that the data packets are arranged in a sorted order corresponding to the identity of the edge module from which the data packet should egress from the distributed data packet switch;~~

~~d) transferring the sorted data packets in fixed length data blocks to a core module of the distributed data packet switch;~~

~~e) switching the fixed length data blocks through the core module to the egress module; and~~

~~f) transferring the payload data packet from the egress module to a subtending sink.~~

New Claims 15-102 added.

15. (New) A packet switch comprising

a plurality of independently-controlled core modules;

a plurality of ingress modules each receiving packets from subtending packet sources and having a link directed to each of said core modules; and

a plurality of egress modules each having a link from each of said core modules and transmitting packets to subtending packet sinks;

wherein each of said ingress modules is operable to:

issue packet-transfer requests each specifying an egress module; and

distribute the packet-transfer requests among said plurality of core modules for scheduling so that each of said core controllers receives a portion of said transfer requests.

16. (New) The packet switch as claimed in claim 15 further including a plurality of core controllers operating concurrently and independently, one core controller associated with each of said independently-controlled core modules and having a packet scheduler.

17. (New) The packet switch as claimed in claim 15 wherein each of said ingress modules is further operable to segment each packet into data segments, each data segment occupying a time slot of a predefined duration.

18. (New) The packet switch as claimed in claim 15 wherein each ingress module has a plurality of ingress ports and an ingress controller, each of the ingress ports having an associated ingress buffer, said ingress controller being adapted to sort packets arriving at the ingress buffers from the subtending packet sources into ingress queues, each ingress queue corresponding to a one of said egress modules.

19. (New) The packet switch as claimed in claim 18 wherein said sorting is a logical sorting.

20. (New) The packet switch as claimed in claim 18 wherein the ingress controller periodically determines ingress-queue occupancy for each respective egress module and sends a capacity-request vector to a selected one of the core controllers, the capacity-request vector including capacity-allocation requests for connections to a subset of the egress modules.

21. (New) The packet switch as claimed in claim 20 wherein each of said core modules computes a schedule in response to receiving a capacity-request vector, the schedule specifying, for each capacity request, time slots in a predefined time frame.

22. (New) The switch as claimed in claim 21 wherein each core module comprises a single-stage space switch.

23. (New) The switch as claimed in claim 22 wherein said space switch is a rotator space switch comprising an input rotator, a plurality of memory devices, and an output rotator.

24. (New) The switch as claimed in claim 21 wherein the core modules, the ingress modules, and the egress modules are spatially distributed.

25. (New) The switch as claimed in claim 24 wherein each core module comprises a number $S \geq 1$ of single-stage space switches, each space switch having a plurality of input ports each connecting to an ingress module and a plurality of output ports each connecting to an egress module.

26. (New) The switch as claimed in claim 25 wherein said single-stage space switch is an optical switch.

27. (New) The switch as claimed in claim 25 wherein each ingress module has a plurality of timing circuits each communicating with a time counter associated with a one of said core modules to realize time coordination between said each ingress module and said one of said core modules.

28. (New) The packet switch as claimed in claim 25 wherein each link carries a wavelength-division-multiplexed optical signal comprising S wavelength channels and wherein each of said single-stage space switches channels of the same wavelength.

29. (New) A method of scheduling performed by a controller of a core module having at least one link from each of a

plurality of ingress modules, said core module comprising a plurality of S ≥1 space switches, each space switch having a plurality of input ports and a plurality of output ports, the method comprising steps of:

creating a data structure comprising:

a first three-dimensional matrix having a space dimension *s* representing space switches associated with the core module, a space dimension *p* representing space-switch input ports, and a time dimension *t* representing time slots in a slotted time frame; and

a second three-dimensional matrix having said space dimension *s*, a space dimension π representing space-switch output ports, and said time dimension *t*;

receiving capacity-allocation requests from the ingress modules, each request specifying an input port *p*, an output port π, and a number K of time slots per time frame; and

selecting a space switch s and a time slot t and, if both entries {s, *p*, t} of the first matrix and {s, π, t} of the second matrix are free, allocating space-switch s and time-slot t and marking entries {s, *p*, t} and {s, π, t} as busy.

30. (New) The method as claimed in claim 29 further including a step of repeating said selecting until at most K time slots are allocated.

31. (New) The method as claimed in claim 29 further including a step of terminating a current connection by setting the value of K to zero.

32. (New) The method as claimed in claim 30 wherein said selecting for each capacity-allocation request considers all time slots in said time frame, then all of said plurality of space switches.

33. (New) The method as claimed in claim 32 further including a step of producing a scheduling matrix that associates each of said input ports with each of said output ports during each of said time slots.

34. (New) The method as claimed in claim 33 wherein said all time slots are considered in a predetermined order.

35. (New) The method as claimed in claim 33 wherein said all of said plurality of switches are considered in a predetermined order.

36. (New) The method as claimed in claim 33 including the further steps of:

assembling said capacity-allocation requests into a capacity-request matrix each entry of which containing a requested capacity allocation; and

attempting to schedule each entry in the matrix to perform core reconfiguration.

37. (New) The method as claimed in claim 36 including the further step of maintaining at each ingress module two scheduling matrices, one in current use and one in update mode.

38. (New) The method as claimed in claim 37 wherein each time a core reconfiguration occurs, a scheduling matrix in use is swapped for a current scheduling matrix.

39. (New) The method as claimed in claim 38 wherein an unused copy of the scheduling matrix is available for update after the core reconfiguration.

40. (New) The method as claimed in claim 39 wherein rows in the scheduling matrix are executed sequentially, one per time slot, until a next core module reconfiguration occurs, and, after core module reconfiguration, processing continues at a next time slot.

41. (New) A distributed packet switch comprising:

a plurality of m cross connectors each having n outer links and n inner links;

a plurality of n core modules each core module comprising a number of space switches;

a plurality of m×n edge modules; and

a plurality of n core controllers, one core controller associated with each of said core modules;

wherein

each of said outer links connects to an edge module and includes Λ channels in each direction to and from the edge module;

each of said inner links connects to a core module and includes Λ channels in each direction to and from the core module; and

said number of space switches per core module does not exceed the ratio Λ/n.

42. (New) The distributed packet switch as claimed in claim 41 further including an edge controller associated with each of said edge modules;

43. (New) The distributed packet switch as claimed in claim 42 wherein said edge modules and said core modules are spatially distributed over a wide geographical area.

44. (New) The distributed packet switch as claimed in claim 43 wherein each of said outer links is a wavelength-division-multiplexed link and each of said inner links is a wavelength-division-multiplexed link.

45. (New) The distributed packet switch as claimed in claim 43 wherein each of said edge modules has means for time coordination with each of said core modules.

46. (New) The distributed packet switch as claimed in claim 45 wherein said means includes a timing circuit.

47. (New) The distributed packet switch as claimed in claim 46 wherein each of said core modules has means for time coordination with each of said edge modules.

48. (New) The distributed packet switch as claimed in claim 47 wherein said means includes a timing circuit.

49. (New) The distributed packet switch as claimed in claim 48 wherein at least one of said space switches in each of said core modules operates in a time-division-multiplexed mode.

50. (New) The distributed packet switch as claimed in claim 49 wherein at least two edge modules transmit time-division-

multiplexed signals to said at least one of said space switches and adjust their transmission times so that the time-division-multiplexed signals are received at said at least one of said space switches in time alignment.

51. (New) The distributed packet switch as claimed in claim 50 wherein one of said edge modules is collocated with a selected one of said core modules and hosts a controller that serves as the core controller of said one of said core modules.

52. (New) The distributed packet switch as claimed in claim 50 wherein each of said edge modules is adapted to transmit capacity-allocation requests to any of said core modules.

53. (New) The distributed packet switch as claimed in claim 52 wherein the core controller associated with said any of said core modules computes a schedule in response to receiving said capacity-allocation requests, the schedule specifying, for each capacity request, time slots in a predefined time frame.

54. (New) The distributed packet switch as claimed in claim 53 wherein at least one of said space switches is an electronic space switch.

55. (New) The distributed packet switch as claimed in claim 53 wherein at least one of said space switches is an optical space switch.

56. (New) The distributed packet switch as claimed in claim 53 wherein at least one of said space switches is an electronic single-stage rotator switch.

57. (New) A packet switch comprising:

a plurality of egress modules, each for transmitting packets on at least one network link;

a plurality of ingress modules, each for receiving packets from at least one network link and capable of making ingress-to-egress-module connection requests for transferring received packets to any other of the egress modules; and

a plurality of core modules, each capable of simultaneously receiving and independently responding to said connection requests from any of the ingress modules, and of providing ingress-to-egress-module connections between any of the ingress modules and any of the egress modules in response to said connection requests.

58. (New) The packet switch as claimed in claim 57 wherein each core module has its own controller for allocating and scheduling resources to the ingress-to-egress-module connections, said own controller operating independently of the other core modules' controllers.

59. (New) The packet switch as claimed in claim 57 wherein each ingress module is connected to each core module by at least one link and each core module is connected to each egress module by at least one link.

60. (New) The packet switch as claimed in claim 59 wherein said resources comprise data paths of sufficient capacities to accommodate said connection requests.

61. (New) The packet switch as claimed in claim 60 wherein said controller of said each core module operates concurrently with the controllers of the other core modules.

62. (New) The packet switch as claimed in claim 61 wherein each core module includes circuitry for time domain multiplexing a plurality of ingress-to-egress-module connections.

63. (New) The packet switch as claimed in claim 62 wherein each of the plurality of ingress-to-egress-module connections starts at the same ingress module and ends at the same egress module.

64. (New) The packet switch as claimed in claim 63 wherein the ingress modules include means for routing a connection through more than one of said core modules.

65. (New) The packet switch as claimed in claim 64 additionally including at least one cross-connector connecting a subset of the ingress modules to the core modules.

66. (New) The packet switch as claimed in claim 65 wherein all of said core modules switch data blocks of the same time duration.

67. (New) The packet switch as claimed in claim 66 wherein each of said core modules maintains its own time reference.

68. (New) The packet switch as claimed in claim 67 wherein each of said ingress modules times the transmission of its data blocks to arrive at a selected one of said core modules at a time determined by said selected one of said core modules.

69. (New) The packet switch as claimed in claim 57 wherein at least one of said ingress modules is provided with a table of preferred paths to each of said egress modules, each path being defined by a core module.

70. (New) The packet switch as claimed in claim 66 wherein at least one of said core modules comprises at least one space switch.

71. (New) The packet switch as claimed in claim 66 wherein at least one of said core modules is a single-stage electronic rotator switch.

72. (New) The packet switch as claimed in claim 66 wherein each ingress module has a plurality of ingress ports and an ingress controller, each of said ingress ports including:

an associated ingress buffer for receiving packets from subtending packet sources; and

means for sorting packets arriving in the ingress buffer into ingress queues each ingress queue corresponding to an egress module from which packets in said each queue are to egress from the switch for delivery to the subtending packet sinks.

73. (New) The packet switch as claimed in claim 72 wherein each ingress module is further adapted to send capacity-allocation requests to a one of said core modules.

74. (New) The packet switch as claimed in claim 73 wherein each ingress module is further adapted to receive connection schedules from each of said core modules.

75. (New) The packet switch as claimed in claim 73 wherein each of said connections is scheduled for transmission over time-slots of equal duration and each of said connection schedules includes time-slot identifiers for each scheduled connection.

76. (New) The packet switch as claimed in claim 73 wherein each ingress module is further adapted to:

create a vector of pointers to the sorted packets; and

assemble said connection schedules into a scheduling matrix so that a non-blank entry in the scheduling matrix indicates an index of the vector of pointers.

77. (New) The packet switch as claimed in claim 76 wherein said ingress controller receives capacity requirements from subtending sources and determines said capacity-allocation requests.

78. (New) The packet switch as claimed in claim 76 wherein the ingress controller periodically determines a number of packets waiting in the ingress queues for each respective egress module and determines said capacity-allocation requests.

79. (New) The packet switch as claimed in claim 75 wherein each core module comprises a plurality of single-stage rotator switches, each rotator switch having a number of input ports collectively adapted to accommodate a number of channels at least equal to the number of ingress modules and a number of output ports collectively adapted to accommodate a number of channels at least equal to the number of egress modules, each ingress module having at least one channel to each of the rotator switches, and each egress module having at least one channel from each of the rotator switches.

80. (New) The packet switch as claimed in claim 79 wherein each ingress module is combined with an egress module to form an integrated edge module.

81. (New) The packet switch as claimed in claim 75 wherein the core modules are co-located at one geographical location.

82. (New) The packet switch as claimed in claim 75 wherein the core modules, the ingress modules, and the egress modules are spatially distributed.

83. (New) The packet switch as claimed in claim 75 wherein each ingress module has a number of timing circuits at least equal to the number of core modules, each of the timing circuits being time-coordinated with a time counter associated with each of said core modules.

84. (New) The packet switch as claimed in claim 83 wherein one edge module is co-located with each core module and said one edge module serves as a controller for the core module.

85. (New) A method of switching a packet through a switch comprising a plurality of ingress modules each having at least one ingress port, a plurality of egress modules each having at least one egress port, and a plurality of core modules, wherein each ingress module is coupled to each core module, each core module is coupled to each egress module, and packets can traverse only one ingress module, one core module and one egress module in moving from an ingress port to an egress port, the method comprising steps of:

receiving a packet from a subtending traffic source at one of the ingress modules;

selecting at the ingress module one of the egress modules to which to send the packet;

sending from the ingress module a connection request to a selected one of the core modules, requesting a connection of a specified capacity to said one of the egress modules; and

at said selected one of the core modules:

determining a feasible capacity allocation,

subtracting said feasible capacity allocation from said specified capacity to update said connection request, and

returning the connection request to said one of the ingress modules.

86. (New) The method as claimed in claim 85 including the further step of said one of the ingress modules sending the connection request to another of the core modules if said feasible capacity allocation is less than said specified capacity.

87. (New) The method as claimed in claim 86 wherein said any one of the core modules is selected at random.

88. (New) The method as claimed in claim 86 wherein said any one of the core modules is selected according to a preferred order.

89. (New) The method as claimed in claim 88 wherein said preferred order is specific to said one of the ingress modules and said one of the egress modules.

90. (New) An ingress module in a packet switch, the ingress module comprising:

an ingress controller;

a plurality of ingress ports each having an ingress buffer for receiving packets from subtending packet sources, each packet indicating a one of predefined destinations;

a plurality of output ports for directing said packets to a plurality of core modules;

means for sorting the packets received in said ingress buffer into ingress queues each corresponding to one of said destinations;

means for storing a set of predefined paths to each of said predefined destinations;

means for formulating connection requests, each of said connection requests specifying a destination and a required capacity allocation; and

means for selecting a candidate path from among said predefined paths for each connection request.

91. (New) The ingress module as claimed in claim 90 wherein said ingress controller performs said sorting, storing, formulating, and selecting.

92. (New) The ingress module as claimed in claim 91 further including means for distributing said connection requests among said predefined paths.

93. (New) The ingress module as claimed in claim 92 further including means for selecting more than one path for a connection.

94. (New) The ingress module as claimed in claim 91 wherein each of said predefined paths is defined by one of said core modules.

95. (New) The ingress module as claimed in claim 94 wherein said core modules are distributed over a wide area.

96. (New) The ingress module as claimed in claim 95 further including a timing circuit adapted to perform time coordination with each of said core modules.

97. (New) A core module in a packet switch, the core module comprising:

at least one space switch having a plurality of input ports and a plurality of output ports;

a core controller adapted to receive connection requests, each connection request specifying a required capacity allocation and a destination selected from among a set of predefined destinations;

means for associating each destination with one of said output ports;

a scheduler associated with said core controller, said scheduler adapted to time the transfer of packets from said input ports to said output ports; and

means for communicating with sources of said connection requests.

98. (New) The core module as claimed in claim 97 wherein said space switch is a bufferless switch operating in a time-multiplexed mode.

99. (New) The core module as claimed in claim 98 wherein said sources are ingress modules each provided with timing circuitry to control the timing of packet transfer.

100. (New) The core module as claimed in claim 99 further including a timing circuit adapted for time coordination with said timing circuitry.

101. (New) The core module as claimed in claim 100 further adapted to receive wavelength-division-multiplexed channels and assign a wavelength channel to each of said at least one space switch.

102. (New) The core module as claimed in claim 101 wherein said core controller determines a feasible capacity allocation for a connection request not exceeding said required capacity allocation and, when the feasible capacity allocation is less than the required capacity allocation, modifies the connection request and returns a modified connection request to a source of the connection request.